




TECH TALK

Issue 150 | December 2024

**Pioneering Tech
Leadership with a
Legacy of Excellence.**



Galaxy Office Automation Pvt. Ltd.

Celebrating Excellence at Galaxy

Galaxy is honored to be recognized as the "**Top Growth TruScale IaaS Partner of The Year**" by **Lenovo** for **FY23-24** in the Asia-Pacific region.

This honour underscores our dedication to delivering seamless integrations for our customers across various industries and highlights our role in fostering resilience in a future driven by AI.

Thank you to our amazing team for their hard work and to Lenovo for this prestigious recognition. Here's to continued success and innovation in the coming years!



Foreword

Dear Readers,

After last month's note, a lot of people asked me about how they should measure the performance of an AI solution. Which performance metrics should be measured? What method should be used for measurement? I will try and answer these questions now.

At the very outset, one should identify the purpose of the AI solution. Is it for improving customer retention, cross selling, fraud detection, automating processes, answering customer queries, or anything else? This should then be further broken up into AI tasks or models being used—e.g. classification, regression, clustering, NLP, recommendation computer visions, etc. Once these have been identified, one should define purpose-based metrics like churn rate, customer lifetime value, retention rate for a customer retention solution, or fraud detection rate and detection latency for a fraud detection solution. Then depending on the AI tasks or models, metrics like accuracy, precision, recall, mean squared error, F1 score, hit rate, etc can be added. Lastly, you should add some metrics to measure the cost of the solution and the business outcomes like cost savings, revenue growth, and customer experience. By regularly monitoring these metrics, the AI solution can be tweaked and tuned to keep improving the scores and consequently the performance of the solution.

At Galaxy Office Automation Pvt. Ltd., our commitment goes beyond implementing AI solutions—we will assist you in defining and measuring the metrics to ensure that your AI solutions deliver continuous ROI.

For those ready to take AI adoption to the next level, our team is here to support you with ongoing performance assessments, training, and optimisation.

Happy reading.



Anoop Pai Dhungat
Chairman & Managing Director



Future is now!

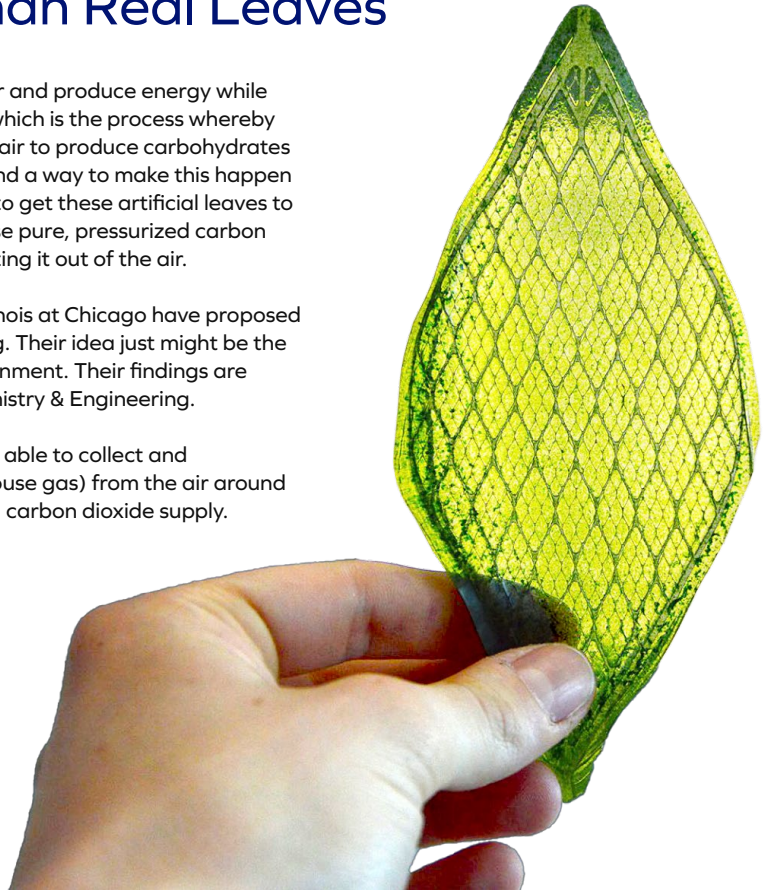



These Artificial Leaves Can Absorb 10 Times More Carbon Dioxide Than Real Leaves

Plants possess a natural ability to purify the air and produce energy while doing so. The ability is called photosynthesis, which is the process whereby plants use water and carbon dioxide from the air to produce carbohydrates using energy from the sun. Scientists have found a way to make this happen artificially. The thing is, they hadn't been able to get these artificial leaves to work outside the lab because the lab leaves use pure, pressurized carbon dioxide from tanks, which is different from getting it out of the air.

But now, researchers from the University of Illinois at Chicago have proposed a design solution that could change everything. Their idea just might be the leaves' ticket out of the lab and into the environment. Their findings are reported in the journal, *ACS Sustainable Chemistry & Engineering*.

The only way that these artificial leaves will be able to collect and concentrate carbon dioxide (a potent greenhouse gas) from the air around us is if they are unhooked from the pressurized carbon dioxide supply.





Here's how Singh and his colleague Aditya Prajapati, a graduate student in his lab, propose to solve this problem:

- The artificial leaf can be placed inside a water-filled capsule constructed out of a semi-permeable membrane.
- When the sunlight warms the water, it evaporates through the membrane, and when that happens the capsule sucks in carbon dioxide.
- The carbon dioxide that's been sucked in then gets converted into carbon monoxide and oxygen by the artificial leaf inside the capsule.
- The carbon monoxide could be siphoned from the device and used to create synthetic fuels ranging from gasoline to methanol.
- The oxygen could be released back into the environment or collected.

In other words, all they have to do is envelope the artificial leaf technology (that has already been developed and works but only in the lab) within this specialised membrane and the whole unit will be able to function outside, like a natural leaf. Furthermore, according to their research, they believe that an artificial leaf built around their design would be 10 times more efficient at converting carbon dioxide to fuel than natural leaves.

Their calculations reveal that 360 of their artificial leaves, each 1.7 meters long and 0.2 meters wide, would generate about half a ton of carbon monoxide daily, which can be used as a basis for synthetic fuels. If those leaves were to be spread out over 500 square meters, then they could reduce the carbon dioxide levels in the air within 100 meters of the space by 10 per cent in just one day.

Read more →

Future Trends and Innovations in Cloud Data Warehousing

Cloud data warehouses help businesses store and manage data in the cloud. They represent a significant evolution in data storage, enabling flexibility, scalability, and affordability in managing increasingly large and complex data.

Every aspect of data management is conducted under the virtual roof of a cloud data warehouse. Most importantly, a cloud data warehouse transforms data into assets companies can use to improve their capabilities, fuel innovation, and enhance profits.

Traditional data warehouses are physical structures typically on-site. While these have served businesses well for many years, a series of challenges, including high costs and complexities with legacy hardware, have rapidly antiquated them. Cloud data warehouses are a robust and ultramodern alternative to traditional data warehouses, one that can lead to profound success for modern businesses. However, enterprises that have to adhere to special compliance or connectivity requirements still leverage on-premises solutions.

By 2026, the market value of the cloud data warehousing industry is forecasted to hit \$12.9 billion, a compound annual growth rate of 22.3%. While North America and Europe hold the highest market share, the fastest-growing segment in cloud data warehousing is the Asia-Pacific region, powered by the booming megamarkets of China and India.



The key factor behind these numbers is that data now drives the world, including business. Cloud data warehouses are highly scalable and provide a safe, secure environment backed up by the expertise of leading high-tech companies.

Industries that benefit the most from cloud data warehousing include manufacturing, energy and utilities, healthcare, IT, government, retail, and BFSI (banking, financial services, and insurance). Since cloud data warehousing is such a flexible solution, the use cases are diverse. But one thing is certain: cloud data warehousing is now the norm and the foundation upon which the future will be constructed.

A Brief Overview of Data Warehouse Architecture

Data warehouse architecture comprises three tiers. The top tier represents the front-end client that offers results via analysis, reporting, data mining tools, and other management. The second or middle tier comprises ETL, which organisations use to access and analyse data. The third or bottom tier of data warehouse architecture is essentially the database server, where enterprises load and store data.

Data can be stored in two ways. Organisations can either leverage high-speed storage to enable quick and frequent access or implement cheap object storage for infrequently accessed data. The data warehouse will move frequently accessed data into "fast" storage to optimise query speeds. Depending on access requirements, organisations can use different logic as well.

Online Analytical Processing (OLAP)

Online analytical processing (OLAP) is a type of data processing that occurs in a data warehouse and serves different workloads and requirements.

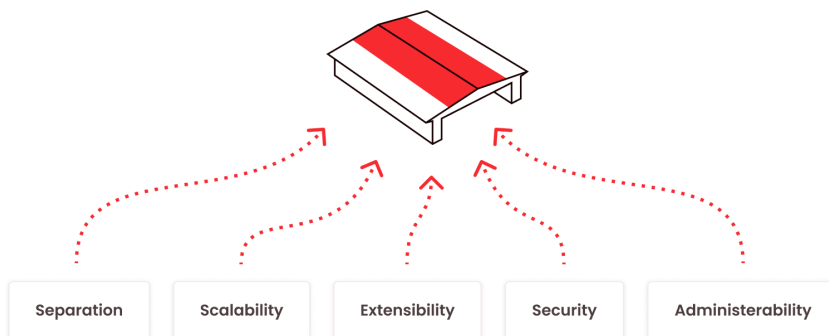
Data consistency is optional for OLAP systems since they typically use data snapshots. OLAP systems handle large data volumes and use denormalised database designs leveraging star schema or snowflake schema. This approach increases data redundancy, improves query performance, and accelerates data-driven decision-making.

How Does Data Warehousing Work?

Data warehouses continuously collect and organise data into a dedicated comprehensive centralised repository. Data collected from various sources are systematically sorted into tables based on the data type and layout.


Insights harvested from a data warehouse help businesses better understand their target audience or customers and be alert to emerging trends. For example, enterprises can gain a competitive advantage by forecasting market changes, formulating a robust pricing strategy, or developing better products.

Properties of data warehouse architecture



There are three layers of data warehousing:

- An enterprise data warehouse (EDW) is a centralised repository providing decision-making support to various departments across the company. EDWs provide a comprehensive and consolidated approach to how companies organise and represent data. As such, data teams can classify the data based on the subject and grant access accordingly.
- An operational data store (ODS) is often a go-to choose for organisations with data warehouse systems failing to satisfy their reporting requirements. As ODS can be refreshed in real-time, it is a popular option for storing routine activities. In a large healthcare system, real-time or near-real-time data access is crucial for patient care. Due to their batch-processing nature, traditional data warehouses can't always meet this need. This is where an ODS comes into play. The ODS can integrate data from these various sources like electronic medical records (EMR), pharmacies, laboratory systems, and radiology to present a unified and current view of the patient's data. For instance, a doctor can access the ODS to get the most recent patient data like lab results or medication history; and a pharmacist can verify prescriptions to prevent harmful drug interactions. Thus, the ODS provides timely, integrated data to healthcare providers, improving patient care.
- A data mart is designed for specific business or industry verticals. For example, they are prevalent in finance, sales, and inventory. Moreover, data marts can quickly collect data from a source.



An EDW stores static data, whereas an ODS integrates dynamic operational data. The data mart creates specialised data views over the EDW.

Enterprises can configure data warehouses into one or multiple of the following system configurations:

- **Offline Operational Database:**
Data will be copied periodically to a server from an ODS to load, process, and report. This approach is practical when data synchronization isn't a must.
- **Offline Data Warehouse:**
Data is stored and regularly updated from the operational database and other sources to derive critical business insights.
- **Real-time Data Warehouse:**
A real-time data warehouse is used for up-to-date insights and analysis based on the latest transactional data. All transactions in an operational database are updated in the data warehouse.
- **Integrated Data Warehouse:**
The integrated data warehouse consolidates data into a unified view for analysis. All transactions occurring in the operational database are simultaneously updated in the data warehouse. Once updated, the data warehouse will generate transactions and forward them to the operational database.

Software tools and hardware used for storing, transforming, and analysing data are called data warehouse appliances.

With the concepts introduced above, we can highlight three key ways in which a data warehouse can work:

- **Basic Data Warehouse:**
Organisations can eliminate data redundancy, which reduces the amount of data in storage. This, in turn, makes data clearer and more user-friendly. The key benefit here is that different departments from multiple sources can quickly access data directly from the warehouse.
- **Data Warehouse with Staging Area:**
Organisations can clean data in "staging areas" before moving it to storage. This is one of the leading methods of ensuring that only relevant and valuable data is stored in the data warehouse.
- **Data Warehouse with Data Marts:**
Organisations can enhance their data warehouse's customisation level after data is processed, allowing them to streamline information to staff, teams, or departments that need it the most. This approach helps boost productivity and accelerate the decision-making pace.

[Read more →](#)

Cloud Repatriation: How to Do It Successfully

The keys to reverse migration success include workload selection, preparing the on-premises infrastructure, and future-proofing the decision to come back in-house.

Cloud repatriation is the process of moving workloads and data back from public cloud infrastructure to on-premises hardware. This could be to a business-owned datacentre, colocation, or other shared facilities.

Organisations might choose to repatriate because of application performance, data security, regulations, or more often cost. Firms will have their own cost-benefit analysis around when to stay with the cloud, or when to move back on-premises, but they also need a plan to make sure any repatriation project is a success.

Why Repatriate

There are no hard-and-fast rules about datasets that benefit most from moving back to on-premises storage. That said, it's possible to identify data, where doing so makes sense. Broadly, repatriation might be the best option where data is sensitive in nature, time-sensitive, or expensive to store in the cloud.

Sensitive data includes regulated information, customer personal data, or where issues of data sovereignty or other regulations put geographical limits on where it can be stored. Governments, too, will have additional restrictions on data that can be stored in the cloud, especially for anything that affects national security.

Time-sensitive data includes information that users need to access as rapidly as possible – think financial trading feeds – or where the application is sensitive to latency. This is a common issue in manufacturing and some areas of R&D, but latency can impact day-to-day business applications, and even technologies such as AI. If an organisation wants complete control over data flows, then it is likely to opt for its own network and storage, not the cloud.

The Cost Factor

Cost, too, is always a factor. Here, it is more a question of how data is used, rather than what it is. It makes a lot of sense to store a long-term archive or backup volume in the cloud, but the calculation changes when organisations want to access the data more frequently. That could be, for example, when using historical data in business intelligence applications or to train AI models. Then, cloud provider egress fees – a charge levied to move data out of the cloud – can mount up. This is one area where the balance between the cloud and on-premises changes over time.

A small test and development server, with minimal storage, will be cost-effective in the cloud, but might be less so if used in production, and carefully calculated cloud storage budgets can be upended if business users decide that data in the cold storage is going to be used on a regular basis instead.

There has been a two-way movement of data and applications for a long time. It's basically a fact of life. People move some things to the cloud because it makes sense, and then after a period, the way they're using that information changes, their needs of it can change, or something else triggers them to modify things, and they move it back.

Preparing Private Infrastructure for Cloud Repatriation

Organisations that want to move data back to their own IT infrastructure, such as a datacentre or Colo facility, need to do the groundwork.

First, they must ensure they have the physical storage capacity for the data being moved. This needs to be planned. Some suppliers have long lead times for new arrays, or even upgrades such as new disks or solid-state modules.

Then there is networking capacity, and physical infrastructure in the datacentre such as rack space, power, and cooling. A large repatriation project might be a prompt to reorganise the datacentre, perhaps by moving to newer equipment that can pack more storage into a single rack or that consumes less power.

Then there are the people needed to support the migration and subsequent day-to-day operations. Are there enough staff to provision and manage a larger system? Do they have the security and privacy skills needed to handle sensitive data? Do they have the technical know-how to handle mission-critical, latency-sensitive applications? These are key questions in a context where many organisations have reduced IT teams, outsourcing to cloud providers.

Enterprises that have grown up in the cloud era might not have the in-house expertise at all. Building up a team can take as long, if not longer, than building up infrastructure, and cost can easily be overlooked while it's wrapped up in cloud service provider fees.

Future-proofing Data and Infrastructure for Repatriation

A key question here is also how to ensure you can reverse the process if you want to. Chief Information Officers (CIOs) will likely want to make sure that if they do move data and applications back from the cloud, they don't miss out on the future benefits of cloud-native applications. In other words, you don't want to move from the cloud to be locked into a local offering forever.

Whether an organisation can maintain its readiness for the full benefits of cloud-native will largely depend on their infrastructure. Use of Kubernetes and other container-based applications on-premises is one way to ensure applications and data are hardware agnostic and easy to port, including to cloud. At the same time, hyperscale cloud suppliers have made it easier to migrate data and provide management tools that can control local and cloud storage.

Nonetheless, the process is rarely simple. There is not easy portability back on-premises, unless you want to use the cloud in a very suboptimal, highly commoditised way. Repatriation can be harder still for firms that use software-as-a-service (SaaS) applications to run business processes.

Lastly, CIOs also need to consider if they can still make use of the cloud for temporary capacity or bursting. Again, this is an area where cloud-native applications and innovations such as object storage and global file systems will help.

Successful repatriation projects will keep a path open to the cloud to support future operations, such as entering a new market where the firm has no datacentre. There, the cloud makes sense, even if the longer-term plan is to bring data back in-house.

NVIDIA Shows AI Model That Can Modify Voices and Generate Novel Sounds

NVIDIA recently showed a new AI model for generating music and audio that can modify voices and generate novel sounds - technology aimed at the producers of music, films, and video games. NVIDIA, the world's biggest supplier of chips and software used to create AI systems, said it does not have immediate plans to publicly release the technology, which it calls Fugatto, short for Foundational Generative Audio Transformer Opus 1. It joins other technologies shown by startups such as Runway and larger players such as Meta Platforms that can generate audio or video from a text prompt.

NVIDIA's version generates sound effects and music from a text description, including novel sounds such as making a trumpet bark like a dog. What makes it different from other AI technologies is its ability to take in and modify existing audio, for example by taking a line played on a piano and transforming it into a line sung by a human voice, or by taking a spoken word recording and changing the accent used and the mood expressed.

"If we think about synthetic audio over the past 50 years, music sounds different now because of computers, because of synthesisers," said Bryan Catanzaro, Vice President of Applied Deep Learning Research at Nvidia. "I think that generative AI is going to bring new capabilities to music, to video games, and to ordinary folks that want to create things."

"While companies such as OpenAI are negotiating with Hollywood studios over whether and how the AI could be used in the entertainment industry, the relationship between tech and Hollywood has become tense, particularly after Hollywood star Scarlett Johansson accused OpenAI of imitating her voice.



Read more →

Google Developing Minimise Option and Other Features for Desktop Mode on Android

Google is speculated to be making efforts to develop its Android operating system (OS) into a more desktop-oriented experience for tablets in a bid to keep up with competitors such as Apple's iPad. It introduced a desktop windowing feature with the Android 15 QPR1 Beta 2 in September which enables use of apps simultaneously in free-form windows. A recent report suggests that the technology giant is developing new functionalities, including an option to minimise windows and launch the same app in multiple windows, for the Desktop Mode on Android.

New Features for Desktop Mode

Android Authority's Mishaal Rahman highlighted the new features in development for the Desktop Mode on Android in a report. The publication suggests that Google is developing a new minimise option which may be added to the header bar. Tapping this option is speculated to minimise the window and it can be maximised by tapping the app's icon in the taskbar — an experience similar to the desktop platform.

The report further highlights that another option for opening a new window could be in development. It was reportedly referenced by Google in its blog post, stating that developers can set a multi-instance property to divulge that the




“system UI should be shown for this app to allow it to be launched as multiple instances.”. This option is speculated to allow users to open a new instance of the app in another window when the desktop windowing feature is toggled.

However, both features are reported to be unavailable in the current builds of Android, but the publication was able to activate them after fiddling with the Android 15 QPR2 Beta 1 release. The report hints at another feature for Chrome which Google detailed in the blog post but isn't live yet — a drag functionality. With this, users may be able to launch a tab in a new window by dragging it. As of now, it is said to be locked behind a Chrome flag.

[Read more →](#)



 A-23/24, Ambika Towers, Ground Floor,
Off. Jijamata Road, Nr. Pump House,
Andheri (E), Mumbai - 400 093, India.

 +91 22 46108999

 marketing@goapl.com

 www.goapl.com